arXiv:1704.04549v1 [math.NA] 14 Apr 2017

Approximate tensor-product preconditioners for very high order discontinuous Galerkin methods

Will Pazner[1] and Per-Olof Persson[2]

[1] *Division of Applied Mathematics, Brown University, Providence, RI, 02912*
[2] *Department of Mathematics, University of California, Berkeley, Berkeley, CA, 94720-3840*

Abstract

In this paper, we develop a new tensor-product based preconditioner for discontinuous Galerkin methods with polynomial degrees higher than those typically employed. This preconditioner uses an automatic, purely algebraic method to approximate the exact block Jacobi preconditioner by Kronecker products of several small, one-dimensional matrices. Traditional matrix-based preconditioners require $\mathcal{O}(p^{2d})$ storage and $\mathcal{O}(p^{3d})$ computational work, where p is the degree of basis polynomials used, and d is the spatial dimension. Our SVD-based tensor-product preconditioner requires $\mathcal{O}(p^{d+1})$ storage, $\mathcal{O}(p^{d+1})$ work in two spatial dimensions, and $\mathcal{O}(p^{d+2})$ work in three spatial dimensions. Combined with a matrix-free Newton-Krylov solver, these preconditioners allow for the solution of DG systems in linear time in p per degree of freedom in 2D, and reduce the computational complexity from $\mathcal{O}(p^9)$ to $\mathcal{O}(p^5)$ in 3D. Numerical results are shown in 2D and 3D for the advection and Euler equations, using polynomials of degree up to $p = 15$. For many test cases, the preconditioner results in similar iteration counts when compared with the exact block Jacobi preconditioner, and performance is significantly improved for high polynomial degrees p.

Keywords: preconditioners; discontinuous Galerkin method; matrix-free

1 Introduction

The discontinuous Galerkin (DG) method, introduced in [22] by Reed and Hill for the neutron transport equation, is a finite element method using discontinuous basis functions. In the 1990s, the DG method was extended to nonlinear systems of conservation laws by Cockburn and Shu [3]. The method has many attractive features, including arbitrarily high formal order of accuracy, and the ability to use general, unstructured meshes with complex geometry. In particular, the promise of a high-order method for fluid flow problems has spurred recent recent interest in the DG method [19]. Higher-order methods promise highly-accurate solutions for less computational cost than traditional low-order methods. Additionally, high-order methods are more computationally intensive per degree of freedom than corresponding low-order methods, resulting in a higher computation-to-communication ratio, and thus rendering these method more amenable to parallelization [1].

High-order accuracy is achieved with the DG method by using a high-degree local polynomial basis on each element in the mesh. There are several challenges that can prevent the use of very high-degree polynomials as basis functions. The number of degrees of freedom per element scales

as $\mathcal{O}(p^d)$, where p is the degree of polynomial approximation, and d is the spatial dimension, resulting in very computationally expensive methods. Using tensor-product evaluations and sum factorizations [18], it is possible to reduce the computational cost of these methods. However, when using explicit time integration schemes, a restrictive stability condition of the form $\Delta t \leq Ch/p^2$, where C is a constant, and h is the element size, requires the use of very small time steps [13]. On the other hand, the DG method couples all the degrees of freedom within each element, so that implicit time integration methods result in block-structured systems of equations, with blocks of size $p^d \times p^d$. Strategies for solving these large linear systems include Newton-Krylov iterative solvers coupled with an appropriate preconditioner [21]. Examples of preconditioners considered include block Jacobi and Gauss-Seidel [17], incomplete LU factorizations (LU) [20], and domain decomposition techniques [5]. Multigrid and multi-level solvers have also been considered [9, 11, 2].

Many of the above preconditioners require the inversion of large the $p^d \times p^d$ blocks corresponding to each element. Using dense linear algebra, this requires $\mathcal{O}(p^{3d})$operations, which quickly becomes intractable. One approach to reduce the computational complexity of implicit methods is to combine Kronecker and sum-factorization techniques with a matrix-free approach. Matrix-free approaches for the DG method have been considered in *e.g.* [4] and [12]. Past work on efficiently preconditioning these systems includes the use of alternating-direction-implicit (ADI) and fast diagonalization method (FDM) preconditioners [6]. Kronecker-product approaches have been studied in the context of spectral methods [24], and applications to the Navier-Stokes equations were considered in [7]. In this work, we describe a new approximate Kronecker-product preconditioner that, when combined with a matrix-free tensor product evaluation approach, allows for efficient solution of the linear systems that arise from implicit time discretizations for high polynomial degree DG methods. This preconditioner requires tensor-product bases on quadrilateral or hexahedral elements. Then, the $p^d \times p^d$ blocks that arise in these systems can be well-approximated by certain Kronecker products of one dimensional $p \times p$ matrices. Using a shuffled singular value decomposition introduced by Van Loan in [26], it is possible to compute decompositions into tensor products of one-dimensional terms that are optimal in the Frobenius norm. Using these techniques, it is possible to construct an approximate tensor-product version of the standard block Jacobi preconditioner, that avoids inverting, or even storing, the large diagonal blocks of the Jacobian matrix.

In Section 2, we give a very brief description of the discontinuous Galerkin method for a general system of hyperbolic conservation laws. In Section 3, we outline the sum-factorization approach, and describe equivalent Kronecker-product representations. Then, in Section 4 we develop the approximate Kronecker-product preconditioners, and provide a new set of algorithms that can be used to efficiently compute and apply these preconditioners. Finally, in Section 5, we apply these preconditioners to several test problems, including the scalar advection equation, and the Euler equations of gas dynamics, in two and three spatial dimensions.

2 Equations and spatial discretization

We give a brief overview of the discontinuous Galerkin method for solving a hyperbolic conservation law of the form

$$\partial_t u + \nabla \cdot F(u) = 0. \tag{1}$$

In order to formulate the method, we first discretize the spatial domain Ω by means of a triangulation $\mathcal{T}_h = \{K_j : \bigcup_j K_j = \Omega\}$. Common choices for the elements K_j of the triangulation are simplex and

block elements. Given a triangulation $\mathcal{T}_h$, we now introduce the finite element space V_h, given by

$$V_h = \left\{ v_h : v_h|_{K_j} \in V(K_j) \right\}, \tag{2}$$

where $V(K_j)$ is a function space local to the element K_j. Such functions admit discontinuities along the element interfaces ∂K_j. In the case of simplex elements, the local function space $V(K_j)$ is taken to be the space of multivariate polynomials of at most degree p, $\mathcal{P}^p(K_j)$. Of particular interest to this paper are the block elements, which in $\mathbb{R}^d$ are defined as the image of the d-fold cartesian product of the interval $[0,1]$ under an isoparametric polynomial transformation map.

By looking for a solution $u_h \in V_h$, multiplying by a test function $v_h \in V_h$, and integrating by parts over each element, we derive the weak formulation of (1),

$$\int_{K_j} (\partial_t u_h) v_h \, dx - \int_{K_j} F(u_h) \cdot \nabla v_h \, dx + \int_{\partial K_j} \widehat{F}(u_h^-, u_h^+, n) v_h \, dA = 0, \qquad \text{for all } K_j \in \mathcal{T}_h, \tag{3}$$

where u_h^- and u_h^+ are the interior and exterior traces (respectively) of u_h on ∂K_j, and $\widehat{F}$ is an appropriately defined *numerical flux function*. The integrals in (3) are approximated using an appropriate quadrature rule, and the resulting system of ordinary differential equations is termed the semi-discrete system. In this work, we use quadrature rules that are given by tensor products of one-dimensional quadratures. Typically, using the method of lines, the time derivative (3) is discretized by means of one of many standard (implicit or explicit) methods for solving ordinary differential equations.

3 The sum-factorization approach

In order to numerically represent the solution, we expand the function u_h in terms of basis functions local to each element. In $\mathbb{R}^d$, the number of degrees of freedom n per element thus scales as $\mathcal{O}(p^d)$. In this work, we will make the assumption that the number of quadrature points, denoted μ, is given by a constant multiple of n, and thus also $\mathcal{O}(p^d)$. We first note that in order to approximate the integrals of the weak form (3), we must evaluate a function $v_h \in V_h$ at each of the quadrature nodes in a given element K. We suppose that $\{\Phi_1, \ldots, \Phi_n\}$ is a basis for the local space $V(K)$. Then, for any $x \in K$, we can expand v_h in terms of its coefficients v_j

$$v_h(x) = \sum_{j=1}^{n} v_j \Phi_j(x). \tag{4}$$

It is important to note that each computation of $v_h(x_\alpha)$ thus requires n evaluations of the basis functions. Performing this computation for each quadrature point therefore requires a total of $\mathcal{O}(p^{2d})$ evaluations. In order to reduce the computational cost of this, and other operations, we describe the sum-factorization approach, first introduced in [18], and extended to the DG method in *e.g.* [27].

3.1 Tensor-product elements

The key to the sum-factorization approach are tensor-product elements, where each element K in the triangulation $\mathcal{T}_h$ is given as the image of the cartesian product $[0,1]^d$ under a transformation

mapping, and where the local basis for each element K is given as the product of one-dimensional basis functions. To be precise, we define the reference element to be the d-dimensional unit cube $\mathcal{R} = [0,1]^d$, and suppose that $T(\mathcal{R}) = K$, where $T : \mathbb{R}^d \to \mathbb{R}^d$ is an isoparametric pth degree polynomial map. We let $\{\phi_j(x)\}_{j=0}^p$ be a basis for $\mathcal{P}^p([0,1])$, the space of polynomials of degree at most p on the unit interval. Then, we define the tensor-product function space

$$V(\mathcal{R}) = \bigotimes_{\ell=1}^{d} \mathcal{P}^p([0,1]) = \operatorname{span}\left\{\phi_{i_1} \otimes \phi_{i_2} \otimes \cdots \otimes \phi_{i_d} : 0 \leq i_k \leq p\right\}, \tag{5}$$

where $\phi_{i_1} \otimes \phi_{i_2} \otimes \cdots \otimes \phi_{i_d}(x_1, x_2, \ldots, x_d) = \phi_{i_1}(x_1)\phi_{i_2}(x_2)\cdots\phi_{i_d}(x_d)$. As a particular example, we consider the one-dimensional nodal basis for $\mathcal{P}^p([0,1])$, with nodes $B = \{b_1, b_2, \ldots, b_{p+1}\} \subseteq [0,1]$, and ϕ_j is the unique degree-p polynomial such that $\phi_j(b_k) = \delta_{jk}$. Thus, the coefficients v_j for a function $v_h \in \mathcal{P}^p([0,1])$ are given by the nodal values $v_j = v_h(b_j)$. The tensor-product basis defined by (5) consists exactly of the multivariate polynomials defined by the nodal basis with nodes given by the d-fold cartesian product $B \times \cdots \times B$. In other words, the basis functions are given by $\Phi_{i_1,i_2,\ldots,i_d}$, which is the unique multivariate polynomial of degree at most p in each variable, such that $\Phi_{i_1,i_2,\ldots,i_d}(b_{j_1}, b_{j_2}, \ldots, b_{j_d}) = \delta_{i_1 j_1}\delta_{i_2 j_2}\cdots\delta_{i_d j_d}$.

For a particular element $K = T(\mathcal{R})$, we define the basis for the space $V(K)$ by means of the transformation map T. Given $x \in K$, we write $x = T(\xi)$, where ξ denotes the *reference coordinate*. Then, we define the basis function $\tilde{\Phi}_{i_1,i_2,\ldots,i_d}(x)$ by $\tilde{\Phi}_{i_1,i_2,\ldots,i_d}(x) = \Phi_{i_1,i_2,\ldots,i_d}(\xi)$. Similarly, it will often be convenient to identify a given function $\tilde{v} \in V(K)$ with the function $v \in V(\mathcal{R})$ obtained by $v = \tilde{v} \circ T^{-1}$.

Given this choice of basis, we also define the quadrature nodes on the d-dimensional unit cube to be the d-fold cartesian product of given one-dimensional quadrature nodes, whose weights are the corresponding products of the one-dimensional weights. Equipped with these choices, we return to the calculation of the quantity (4). For the sake of concreteness, we consider the case where $d = 3$, for which the calculation above should naively require $\mathcal{O}(p^6)$ evaluations. We suppose that the one-dimensional quadrature points are given as $x_1, x_2, \ldots, x_\mu$, and hence we can write the three-dimensional quadrature points as $\boldsymbol{x}_{\alpha,\beta,\gamma} = (x_\alpha, x_\beta, x_\gamma)$, for all $1 \leq \alpha, \beta, \gamma \leq \mu$. We then factor the summation in (4) to obtain

$$\begin{aligned} v_h(\boldsymbol{x}_{\alpha,\beta,\gamma}) &= \sum_{i,j,k=1}^{p+1} v_{ijk}\Phi_{ijk}(x_\alpha, x_\beta, x_\gamma) \\ &= \sum_{k=1}^{p+1}\sum_{j=1}^{p+1}\sum_{i=1}^{p+1} v_{ijk}\phi_i(x_\alpha)\phi_j(x_\beta)\phi_k(x_\gamma) \\ &= \sum_{k=1}^{p+1}\phi_k(x_\gamma)\sum_{j=1}^{p+1}\phi_j(x_\beta)\sum_{i=1}^{p+1} v_{ijk}\phi_i(x_\alpha). \end{aligned} \tag{6}$$

We notice that the index of each summation ranges over $p+1$ values, and there are three free indices in each sum. Thus, the total number of operations required to evaluate a function v_h at each of the quadrature points is $\mathcal{O}(p^4)$. For general dimension d, this computation requires $\mathcal{O}(p^{d+1})$ basis function evaluations. Thus, the computational work *per degree of freedom* is linear in the degree p of polynomial basis, in contrast to the original estimate of $\mathcal{O}(p^d)$ computational work per degree of freedom, which is exponential in spatial dimension d. In a similar fashion, the tensor-product

structure of this function space can be exploited in order to compute the integrals in equation (3) in linear time per degree of freedom.

3.2 Kronecker-product structure

The sum-factorization procedure shown in (6) can be described simply and elegantly as a linear-algebraic Kronecker product structure. We recall that the Kronecker product of matrix $A^{k\times\ell}$ and $B^{m\times n}$ (whose dimensions are indicated by the superscripts), is the $km \times \ell n$ matrix C defined by

$$A \otimes B = C = \begin{pmatrix} a_{11}B & a_{12}B & \cdots & a_{1\ell}B \\ a_{21}B & a_{22}B & \cdots & a_{2\ell}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1}B & a_{k2}B & \cdots & a_{k\ell}B \end{pmatrix}. \tag{7}$$

The Kronecker product has many desirable and useful properties, enumerated in Van Loan's exposition [15].

We can define the one-dimensional *Gauss point evaluation matrix* as the $\mu\times(p+1)$ Vandermonde-type matrix obtained by evaluating each of the one-dimensional basis functions at all of the quadrature points,

$$G_{\alpha j} = \phi_j(x_\alpha). \tag{8}$$

Then, in the three-dimensional case described above, we can define a vector v of length $(p+1)^3$ whose entries are given by the entries of the 3-tensor v_{ijk}, concatenated. Thus, we obtain the values of v_h evaluated at the quadrature points $\boldsymbol{x}_{\alpha,\beta,\gamma}$ by computing

$$v_h(\boldsymbol{x}_{\alpha,\beta,\gamma}) = (G \otimes G \otimes G)v. \tag{9}$$

This Kronecker-product representation is computationally equivalent to the sum-factorized version from the preceding section. Indeed, many of the operations needed for the computation of the discontinuous Galerkin method are amenable to being written in the form of Kronecker products.

For instance, it is often useful to approximate quantities of the form

$$\int_K f(x)v_h(x)\ dx, \tag{10}$$

where f is an arbitrary function whose value is known at the appropriate quadrature nodes. This requires approximating the integrals

$$\int_K f(x)\tilde{\Phi}_{ijk}(x)\ dx \tag{11}$$

for each of the basis functions $\tilde{\Phi}_{ijk}$. We consider the element K to be the image under the isoparametric transformation map T of the reference element $\mathcal{R} = [0,1]^d$. In this notation, for all $x \in K$, $x = T(\xi)$, where $\xi \in \mathcal{R}$. We then write (11) as an integral over the reference element,

$$\int_K f(x)\tilde{\Phi}_{ijk}(x)\ dx = \int_{\mathcal{R}} f(T(\xi))\Phi_{ijk}(\xi)|\det(T'(\xi))|\ d\xi. \tag{12}$$

To this end, we define a diagonal weight matrix W by whose entries along the diagonal are given by w_α, where w_α is the quadrature weight associated with the point x_α. Additionally, we define the

Operation	2D	3D
Evaluate solution at quadrature points	$(G \otimes G)\, u$	$(G \otimes G \otimes G)\, u$
Integrate function f (known at quadrature points) against test functions	$\left(G^T W \otimes G^T W\right) J_T f$	$\left(G^T W \otimes G^T W \otimes G^T W\right) J_T f$
Integrate function $f = (f_1, \ldots, f_d)$ against gradient of test functions	$\left(G^T W \otimes D^T W\right) J_T f_1$ $\left(D^T W \otimes G^T W\right) J_T f_2$	$\left(G^T W \otimes G^T W \otimes D^T W\right) J_T f_1$ $\left(G^T W \otimes D^T W \otimes G^T W\right) J_T f_2$ $\left(D^T W \otimes G^T W \otimes G^T W\right) J_T f_3$

Table 1: Kronecker-product form of DG operations

$\mu^3 \times \mu^3$ diagonal matrix J_T whose entries are equal to the Jacobian determinant of the isoparametric mapping $\det(T'(\xi))$ at each of the quadrature points. Then, the $(p+1)^3$ integrals of the form (12) can be found as the entries of the vector

$$\left(G^T W\right) \otimes \left(G^T W\right) \otimes \left(G^T W\right) J_T f(\boldsymbol{x}_{\alpha,\beta,\gamma}). \tag{13}$$

In a similar fashion, the computation of all of the quantities needed to formulate a DG method can be written in Kronecker form. In Table 1, we briefly summarize the Kronecker-product formulation of several other important operations needed for the DG method.

3.3 Explicit time integration

It is important to note that all of these operations have a computational complexity of at most $\mathcal{O}(p^{d+1})$. In other words, the cost of these operations scales linearly in p per degree of freedom. We now return to the semi-discrete system of equations (3), which we rewrite as

$$M(\partial_t u_h) = r, \tag{14}$$

where r is the quadrature approximation of the second two integrals on the left-hand side of (3). Then, using the operations described above, it is possible to compute all the integrals required to form r. What is required in order to integrate this semi-discrete equation explicitly in time is to invert the mass matrix M. We recall that, with a tensor-product basis, we can compute the element-wise mass matrix (on, *e.g.*, element K_j) as

$$M_j = \left((G^T W) \otimes \cdots \otimes (G^T W)\right) J_T \left(G \otimes \cdots \otimes G\right) \tag{15}$$

where G is the Gauss point evaluation matrix defined above, W is the diagonal matrix with the one-dimensional quadrature weights on the diagonal, and J_T is the $(\mu^d) \times (\mu^d)$ diagonal matrix whose entries are equal to the absolute Jacobian determinant of the element transformation map, evaluated at each of the quadrature points.

One strategy, proposed in [14], is to use the same number of quadrature points as DG nodes, such that $\mu = p + 1$. In that case, all of the matrices appearing in (15) are square, and we can compute

$$M_j^{-1} = \left(G^{-1} \otimes \cdots \otimes G^{-1}\right) J_T^{-1} \left((G^T W)^{-1} \otimes \cdots \otimes (G^T W)^{-1}\right). \tag{16}$$

Since G and $G^T W$ are both $(p+1) \times (p+1)$ matrices, these operations can be performed in $\mathcal{O}(p^3)$ time. Additionally, J_T is a $(p+1)^d \times (p+1)^d$ diagonal matrix, and thus can be inverted in $(p+1)^d$ operations. (On a practical note, in this case we would avoid explicitly forming the inverse matrices G^{-1} and $(G^T W)^{-1}$, and would instead opt to form their LU factorizations). Thus, the linear system (14) can be solved in the same complexity as multiplying by the expression on the right-hand side of (16), *i.e.* $\mathcal{O}(p^{d+1})$.

A serious drawback to this approach is that using the same number of quadrature points as DG nodes does not, in general, allow for exact integration of the quantity

$$\int_K u_h v_h \, dx, \qquad u_h, v_h \in V_h, \tag{17}$$

because of the use of isoparametric elements, where the Jacobian determinant of the transformation mapping may itself be a high-degree polynomial. In order to address this issue, we introduce a new strategy for solving the system (14). We first note that the global mass matrix has a natural element-wise block-diagonal structure, where, furthermore, each block M_j is a symmetric positive-definite matrix. Thus, we can solve this system of equations element-by-element, using the preconditioned conjugate gradient (PCG) method. As a preconditioner, we use the under-integration method described above. This has the consequence that if the element transformation mapping is bilinear (and so the corresponding element is straight-sided) then its Jacobian determinant is linear, and with the appropriate choice of quadrature points, the integral (17) can be computed exactly, and the PCG method will converge within one iteration. In practice, we observe that the number of PCG iterations required to converge is very small even on curved, isoparametric meshes, and does not grow with p.

The techniques described above are sufficient to implement an explicit discontinuous Galerkin method with tensor-product elements, requiring $\mathcal{O}(p^{d+1})$ operations per time step. The main restriction to using such explicit methods with very high polynomial degree p is the restrictive CFL condition, that requires that the time step satisfy approximately $\Delta t \leq Ch/p^2$ [13], which can prove to be prohibitively expensive. For this reason, we are interested in applying some of the same tensor-product techniques to efficiently integrate in time implicitly, and thus avoid the overly-restrictive CFL condition.

3.4 Implicit time integration

Instead of an explicit time integration method, we now consider an implicit schemes such as backward differentiation formulas (BDF) or diagonally-implicit Runge-Kutta (DIRK) methods. The main advantage of such methods is that they remain stable for larger time steps, even in the presence of highly anisotropic elements. Additionally, these methods avoid the restrictive p-dependent explicit stability condition mentioned above. Such implicit methods require the solution of systems of the form

$$M u_h - \Delta t f(u_h) = r, \tag{18}$$

which, when solved by means of Newton's method, give rise to linear systems of the form

$$(M - \Delta t J) x = b, \tag{19}$$

where the matrix J is the Jacobian of the potentially non-linear function f.

The most immediate challenge towards efficiently implementing an implicit method for high polynomial degree on tensor product elements is forming the Jacobian matrix. In general, all the

degrees of freedom within one element are coupled, and thus the diagonal blocks of the Jacobian matrix corresponding to a single element are dense $(p+1)^d \times (p+1)^d$ matrices. Therefore, it is impossible to explicitly form this matrix in less than $\mathcal{O}(p^{2d})$ time. To circumvent this, using the techniques described in the preceding sections, it is possible to solve the linear systems arising from implicit time integration by means of an iterative method such as GMRES [21]. Each iteration requires performing a matrix-vector multiplication by the mass matrix and the Jacobian matrix. If we avoid explicitly forming these matrices, then the multiplications can be performed in $\mathcal{O}(p^{d+1})$ time, using methods similar to those described in the preceding section.

3.4.1 Matrix-free tensor-product Jacobians

In order to efficiently implement the implicit method described above, we first apply the sum-factorization technique to efficiently evaluate the matrix-vector product

$$(M - \Delta t J)v \tag{20}$$

for a given vector v. As described in Section 3.2, the mass matrix is a block diagonal matrix whose jth block can be written in Kronecker form as

$$M_j = \left((G^T W) \otimes \cdots \otimes (G^T W)\right) J_T \left(G \otimes \cdots \otimes G\right), \tag{21}$$

where G is the $\mu \times (p+1)$ basis evaluation matrix, J_T is the diagonal transformation Jacobian matrix, and W is the diagonal quadrature weight matrix. Multiplying by the Kronecker products can be performed with $\mathcal{O}(p^{d+1})$ operations, and multiplying by J_T requires exactly $(p+1)^d$ operations. Thus, the product Mx requires $\mathcal{O}(p^{d+1})$ operations. We now describe our algorithm for also computing the product Jx with the same complexity. For simplicity of presentation, we will take $d = 2$, but the algorithm is immediately generalizable to arbitrary dimension d.

We first consider an element-wise blocking of the matrix J. Each block is a $(p+1)^2 \times (p+1)^2$ matrix, with blocks along the diagonal corresponding to each element in the triangulation, and blocks off the diagonal corresponding to the coupling between neighboring elements through their common face. We write $r = f(u)$, and restrict our attention to one element. We define the indices $1 \le i, j, k, \ell \le p+1$, such that the entries of the diagonal block of the Jacobian can be written as

$$\frac{\partial r_{ij}}{\partial u_{k\ell}}. \tag{22}$$

We define r_{ij} by

$$r_{ij} = \int_K F(u) \cdot \nabla \tilde{\Phi}_{ij} \, dx - \int_{\partial K} \widehat{F}(u^-, u^+, n) \tilde{\Phi}_{ij} \, dA \tag{23}$$

which we evaluate using the quadrature rule

$$\begin{aligned} r_{ij} = &\sum_{\alpha=1}^{\mu} \sum_{\beta=1}^{\mu} w_\alpha w_\beta F(u(x_\alpha, x_\beta)) \cdot \nabla \left(\phi_i(x_\alpha)\phi_j(x_\beta)\right) \\ &- \sum_{e \in \partial K} \sum_{\alpha=1}^{\mu} w_\alpha \widehat{F}(u^-(x_\alpha^e, y_\alpha^e), u^+(x_\alpha^e, y_\alpha^e), n(x_\alpha^e, y_\alpha^e))\phi_i(x_\alpha^e)\phi_j(y_\alpha^e), \end{aligned} \tag{24}$$

where the notation (x_α^e, y_α^e) represents the coordinates of the αth quadrature node alone the face e of ∂K. We also recall that the function u is evaluated by expanding in terms of the local basis functions, *e.g.*

$$u(x_\alpha, x_\beta) = \sum_{k=1}^{p+1}\sum_{\ell=1}^{p+1} u_{k\ell}\phi_k(x_\alpha)\phi_\ell(x_\beta), \tag{25}$$

which can be evaluated efficiently as

$$(G \otimes G)\, u. \tag{26}$$

Thus, the entries of the Jacobian can be written as

$$\begin{aligned}\frac{\partial r_{ij}}{\partial u_{k\ell}} &= \sum_{\alpha=1}^{\mu}\sum_{\beta=1}^{\mu} w_\alpha w_\beta \phi_k(x_\alpha)\phi_\ell(x_\beta)\frac{\partial F}{\partial u}(u(x_\alpha, x_\beta)) \cdot \nabla\left(\phi_i(x_\alpha)\phi_j(x_\beta)\right) \\ &\quad - \sum_{e\in\partial K}\sum_{\alpha=1}^{\mu} w_\alpha \phi_k(x_\alpha^e)\phi_\ell(y_\alpha^e)\frac{\partial \widehat{F}}{\partial u^-}(u^-(x_\alpha^e, y_\alpha^e), u^+(x_\alpha^e, y_\alpha^e), n(x_\alpha^e, y_\alpha^e))\phi_i(x_\alpha^e)\phi_j(y_\alpha^e).\end{aligned} \tag{27}$$

Since there are $(p+1)^{2d}$ such entries, we avoid explicitly computing the entries of this matrix, and instead describe how to compute the matrix-vector product Jx. As a pre-computation step, we compute the flux derivatives $\partial F/\partial u$ and numerical flux derivatives $\partial \widehat{F}/\partial u^-$ at each of the quadrature nodes (x_α, x_β) in the element K. For simplicity, we introduce the notation

$$\frac{\partial F}{\partial u}(x_\alpha, x_\beta) = \frac{\partial F}{\partial u}(u(x_\alpha, x_\beta)), \qquad \frac{\partial \widehat{F}}{\partial u^-}(x_\alpha^e, y_\alpha^e) = \frac{\partial \widehat{F}}{\partial u^-}(u^-(x_\alpha^e, y_\alpha^e), u^+(x_\alpha^e, y_\alpha^e), n(x_\alpha^e, y_\alpha^e)). \tag{28}$$

Then, applying the sum-factorization technique, the terms of the product of Jv for a given vector v corresponding to the diagonal block takes the form

$$\begin{aligned}\left(\frac{\partial r_{ij}}{\partial u_{k\ell}}\right) v_{k\ell} &= \sum_{k=1}^{p+1}\sum_{\ell=1}^{p+1}\sum_{\alpha=1}^{\mu}\sum_{\beta=1}^{\mu} w_\alpha w_\beta \phi_k(x_\alpha)\phi_\ell(x_\beta)\frac{\partial F}{\partial u}(x_\alpha, x_\beta) \cdot \nabla\left(\phi_i(x_\alpha)\phi_j(x_\beta)\right) v_{k\ell} \\ &\quad - \sum_{k=1}^{p+1}\sum_{\ell=1}^{p+1}\sum_{e\in\partial K}\sum_{\alpha=1}^{\mu} w_\alpha \phi_k(x_\alpha^e)\phi_\ell(y_\alpha^e)\frac{\partial \widehat{F}}{\partial u^-}(x_\alpha^e, y_\alpha^e)\phi_i(x_\alpha^e)\phi_j(y_\alpha^e)v_{k\ell}\end{aligned} \tag{29}$$

$$\begin{aligned}&= \sum_{\alpha=1}^{\mu} w_\alpha \phi_i'(x_\alpha)\sum_{\beta=1}^{\mu} w_\beta \frac{\partial F_1}{\partial u}(x_\alpha, x_\beta)\phi_j(x_\beta)\sum_{\ell=1}^{p+1}\phi_\ell(x_\beta)\sum_{k=1}^{p+1}\phi_k(x_\alpha)v_{k\ell} \\ &\quad + \sum_{\alpha=1}^{\mu} w_\alpha \phi_i(x_\alpha)\sum_{\beta=1}^{\mu} w_\beta \frac{\partial F_2}{\partial u}(x_\alpha, x_\beta)\phi_j'(x_\beta)\sum_{\ell=1}^{p+1}\phi_\ell(x_\beta)\sum_{k=1}^{p+1}\phi_k(x_\alpha)v_{k\ell} \\ &\quad + \sum_{e\in\partial K}\sum_{\alpha=1}^{\mu} w_\alpha \frac{\partial \widehat{F}}{\partial u^-}(x_\alpha^e, y_\alpha^e)\phi_i(x_\alpha^e)\phi_j(y_\alpha^e)\sum_{\ell=1}^{p+1}\phi_\ell(y_\alpha^e)\sum_{k=1}^{p+1}\phi_k(x_\alpha^e)v_{k\ell},\end{aligned} \tag{30}$$

where F_1 and F_2 are the x and y components of the flux function F, respectively. We notice that in each of the above summations, there are at most two free indices, and therefore each sum can be computed in $\mathcal{O}(p^3)$ time, achieving linear time in p per degree of freedom. The terms of the

product corresponding to the off-diagonal blocks have a similar form to the face integral in the above equations, and can similarly be computed in $\mathcal{O}(p^3)$ time.

To summarize, we describe the algorithm for computing the matrix-products of the form Jx in Algorithm 1.

Algorithm 1 Matrix-free computation of Jv

1: **Pre-computation:**
2: Evaluate the solution at quadrature points by computing $(G \otimes G)\, u$ ▷ *Complexity:* $\mathcal{O}(p^3)$
3: Evaluate the flux Jacobians $\frac{\partial F}{\partial u}$ and $\frac{\partial \widehat{F}}{\partial u^-}$ at quadrature points ▷ *Complexity:* $\mathcal{O}(p^2)$
4: Compute the matrix-vector product using the sum-factorized form ▷ *Complexity:* $\mathcal{O}(p^3)$

The first two operations can be performed as a pre-computation step, and only the third step need be repeated when successively multiplying the same Jacobian matrix by different vectors (as in the case of an iterative linear solver).

4 Tensor-product preconditioners

One of the main challenges in successfully applying such a matrix-free method is preconditioning [23]. Common preconditioners typically used for implicit DG methods include block Jacobi, block Gauss-Seidel, and block ILU preconditioners [21]. Computing these preconditioners first requires forming the matrix, and additionally requires the inversion of certain blocks. Typically, this would incur a cost of $\mathcal{O}(p^{3d})$, which quickly grows to be prohibitive as we take p to be large. To remedy this issue, we develop a preconditioner for two and three spatial dimensions that takes a similar Kronecker product form to those seen in the previous section.

We draw inspiration from the tensor-product structure often seen in finite-difference and spectral approximations to, *e.g.* the Laplacian operator on a n^d cartesian grid, which can be written in one, two, and three spatial dimensions, respectively, as

$$L_{1\mathrm{D}} = T_n, \qquad L_{2\mathrm{D}} = I \otimes T_n + T_n \otimes I, \qquad L_{3\mathrm{D}} = I \otimes I \otimes T_n + I \otimes T_n \otimes I + T_n \otimes I \otimes I, \tag{31}$$

where T_n is the standard one-dimensional approximation to the Laplacian. Given a general conservation law of the form (1), the flux function F is not required to possess any particular structure, and thus the DG discretization of such a function will not be exactly expressible in a similar tensor-product form. That being said, many of the key operations in DG, listed in Table 1, are expressible in a similar form. Therefore, in order to precondition the implicit systems of the form

$$(M - \Delta t J)u = b, \tag{32}$$

we look for tensor-product approximations to the diagonal blocks A of the matrix $M - \Delta t J$. Specifically, in two spatial dimensions, we are interested in finding preconditioners P of the form

$$A \approx P = \sum_{j=1}^{r} A_j \otimes B_j, \tag{33}$$

for a fixed number of terms r, where each of the matrices A_j and B_j are of size $(p+1) \times (p+1)$. Similarly, in three dimensions, we look for P which is given as the sum of three-term Kronecker

products. Given r, it is possible to find the best possible approximation (in the Frobenius norm) of the form (33) to an arbitrary given matrix by means of the Kronecker-product singular value decomposition (KSVD).

4.1 Kronecker-product singular value decomposition

In [26], Van Loan posed the *nearest Kronecker product problem* (NKP): given a matrix $A \in \mathbb{R}^{m\times n}$ (with $m = m_1 m_2$ and $n = n_1 n_2$), and given a fixed number r, find matrices $A_j \in \mathbb{R}^{m_1\times n_1}, B_j \in \mathbb{R}^{m_2\times n_2}$ that minimize the Frobenius norm

$$\left\| A - \sum_{j=1}^{r} A_j \otimes B_j \right\|_F . \tag{34}$$

The solution to the NKP given by Van Loan is as follows. We first consider the "blocking" of A:

$$A = \begin{pmatrix} A_{11} & A_{12} & \cdots & A_{1,n_1} \\ A_{21} & A_{22} & \cdots & A_{2,n_1} \\ \vdots & \vdots & \ddots & \vdots \\ A_{m_1,1} & A_{m_1,2} & \cdots & A_{m_1,n_1} \end{pmatrix}, \tag{35}$$

where each block is a $m_2 \times n_2$ matrix. We then define a rearranged (or *shuffled*) version $\widetilde{A}$ of the matrix A, which is a $m_1 n_1 \times m_2 n_2$ matrix given by

$$\widetilde{A} = \begin{pmatrix} \widetilde{A}_1 \\ \widetilde{A}_2 \\ \vdots \\ \widetilde{A}_{n_1} \end{pmatrix}, \text{ where } \widetilde{A}_j \text{ is a block of rows given by } \widetilde{A}_j = \begin{pmatrix} \mathrm{vec}(A_{1j})^T \\ \mathrm{vec}(A_{2j})^T \\ \vdots \\ \mathrm{vec}(A_{m_1,j})^T \end{pmatrix}, \tag{36}$$

where the vec operator is defined so that $\mathrm{vec}(A_{ij})$ is the column vector of length $m_2 n_2$ obtained by "stacking" the columns of A_{ij}. This rearranged matrix has the property that, given matrices A_j, B_j

$$\left\| A - \sum_{j=1}^{r} A_j \otimes B_j \right\|_F = \left\| \widetilde{A} - \sum_{j=1}^{r} \mathrm{vec}(A_j)\,\mathrm{vec}(B_j)^T \right\|_F , \tag{37}$$

and therefore the NKP problem (34) has been reduced the finding the closest rank-r approximation to $\widetilde{A}$. This approximation can be found by computing the singular value decomposition (SVD) of $\widetilde{A}$,

$$\widetilde{A} = U\Sigma V^T, \tag{38}$$

then the solution to (34) is given by reshaping the columns of U and V, such that

$$\mathrm{vec}(A_j) = \sqrt{\sigma_j} U_j, \qquad \mathrm{vec}(B_j) = \sqrt{\sigma_j} V_j. \tag{39}$$

This construction is referred to as the *Kronecker product SVD* (KSVD).

4.1.1 Efficient computation of the KSVD

In general, computing the singular value decomposition of a matrix is an expensive process, with cubic complexity. However, if the number r of desired terms in the summation (33) is much smaller than the rank $\tilde{r}$ of the matrix $\widetilde{A}$, then it is possible to well-approximate the largest singular values and associated left and right singular vectors by means of a Lanczos algorithm [8]. This algorithm has the additional advantage that an explicit representation of the matrix $\widetilde{A}$ is not required, rather only the ability to multiply vectors by the shuffled matrices $\widetilde{A}$ and $\widetilde{A}^T$. In this section, we follow the presentation from [26]. The Lanczos SVD procedure is described in Algorithm 2.

Algorithm 2 Lanczos SVD

1: $v_0 \leftarrow$ random vector with $\|v_0\|_2 = 1$
2: $p_0 \leftarrow v_0, \beta_0 \leftarrow 1, u_0 \leftarrow 0$
3: **for** $j = 0$ to J (maximum number of iterations) **do**
4: $\quad v_j \leftarrow p_j / \beta_j$
5: $\quad r_j \leftarrow \widetilde{A} v_j - \beta_j u_j$
6: $\quad$ Orthogonalize r_j.
7: $\quad \alpha_j \leftarrow \|r_j\|_2$
8: $\quad u_{j+1} \leftarrow r_j / \alpha_j$
9: $\quad p_{j+1} \leftarrow \widetilde{A}^T u_j - \alpha_j v_j$
10: $\quad$ Orthogonalize p_{j+1}.
11: $\quad \beta_{j+1} \leftarrow \|p_{j+1}\|_2$
12: $\quad$ **if** $|\beta_{j+1}| <$ tolerance **then**
13: $\quad\quad$ **break**
14: $U \leftarrow (u_1, u_2, \ldots, u_{j+1})$
15: $V \leftarrow (v_0, v_1, \ldots, v_j)$
16: Construct bidiagonal matrix B, with diagonal $\alpha_0, \ldots, \alpha_j$, and superdiagonal $\beta_1, \ldots, \beta_j$.
17: Compute r largest singular values σ_k (and corresponding left and right singular vectors, u'_k, v'_k) of B
18: Singular values of $\widetilde{A}$ are σ_j, singular vectors are $U u'_k$ and $V v'_k$.

We remark that there are many variations on the orthogonalization procedure referred to in lines 6 and 10 of Algorithm 2, including partial or full orthogonalization. In this work, we perform full orthogonalization of the vectors u_j and v_j at each iteration of the Lanczos algorithm.

As mentioned previously, one of the key advantages of the Lanczos algorithm is that an explicit representation of the matrix $\widetilde{A}$ can be foregone, since only matrix-vector products of the form $\widetilde{A}x$ and $\widetilde{A}^T x$ are required. As described in [26], we can compute these matrix-vector products according to Algorithms 3 and 4. Taking advantage of the specific tensor-product form of the matrix A, and using techniques similar to those used for the matrix-free Jacobian evaluation from Section 3.4.1, it is possible to efficiently compute the matrix-vector products. Specialized kernels are required for two and three spatial dimensions, and the details of this process are described in the following sections.

Algorithm 3 Compute $u = \widetilde{A}v$

$u \leftarrow 0$
for $i = 1$ to n_1 **do**
    rows $\leftarrow (i-1)m_1 + 1, \ldots, im_1$
    for $j = 1$ to n_2 **do**
        Define $Z \in \mathbb{R}^{m_2 \times m_1}$ by $\text{vec}(Z) = A(:, (i-1)n_2 + j)$
        $u(\text{rows}) \leftarrow u(\text{rows}) + Z^T v((j-1)m_2 + 1{:}jm_2)$

Algorithm 4 Compute $u = \widetilde{A}^T v$

$u \leftarrow 0$
for $i = 1$ to n_2 **do**
    rows $\leftarrow (i-1)m_2 + 1, \ldots, im_2$
    for $j = 1$ to n_1 **do**
        Define $Z \in \mathbb{R}^{m_2 \times m_1}$ by $\text{vec}(Z) = A(:, (j-1)n_2 + i)$
        $u(\text{rows}) \leftarrow u(\text{rows}) + Zv((j-1)m_1 + 1{:}jm_1)$

4.2 Two spatial dimensions

Having shown that, given the number of terms r in the sum, it is possible to find the best approximation of the form (33), we now address the issue of solving linear systems of equations with such a matrix. In the case that $r = 1$, we have $P = A_1 \otimes B_1$, and it is clear that $P^{-1} = A_1^{-1} \otimes B_1^{-1}$, and thus the $(p+1)^2 \times (p+1)^2$ problem is reduced to two problems of size $(p+1) \times (p+1)$. Our experience has shown that $r = 1$ is not sufficient to accurately approximate the Jacobian matrix, and the resulting preconditioners are not very effective. For this reason, we choose $r = 2$, and obtain a linear system of the form

$$Px = (A_1 \otimes B_1 + A_2 \otimes B_2)x = b. \tag{40}$$

Because of the additional term in this sum, it is not possible to invert this matrix factor-wise. Instead, we follow the matrix diagonalization technique described in [24, 16]. We multiply the system of equations on the left by $(A_2^{-1} \otimes B_1^{-1})$ to obtain

$$(A_2^{-1}A_1 \otimes I + I \otimes B_1^{-1}B_2)x = (A_2^{-1} \otimes B_1^{-1})b. \tag{41}$$

We let $C_1 = A_2^{-1}A_1$ and $C_2 = B_1^{-1}B_2$. We then remark that if C_1 and C_2 are diagonalizable matrices, the sum $C_1 \otimes I + I \otimes C_2$ can be simultaneously diagonalized by means of the eigendecomposition. More generally, the Schur factorization of the matrices C_1 and C_2 is guaranteed to exist, and thus the summation $C_1 \otimes I + I \otimes C_2$ can be simultaneously (quasi)-triangularized by of the (real) Schur decomposition. That is to say, we find orthogonal transformation matrices Q_1 and Q_2 such that

$$C_1 = Q_1 T_1 Q_1^T, \tag{42}$$
$$C_2 = Q_2 T_2 Q_2^T, \tag{43}$$

where T_1 and T_2 are quasi-triangular matrices. Our numerical experiments have indicated that the Schur factorizations results in better numerical conditioning than the eigendecomposition, and thus

we elect to triangularlize the matrix rather than diagonalize. Therefore, we can reformulate the linear system as

$$\begin{aligned}(Q_1 \otimes Q_2)(T_1 \otimes I + I \otimes T_2)(Q_1^T \otimes Q_2^T)x &= \left(Q_1 T_1 Q_1^T \otimes I + I \otimes Q_2 T_2 Q_2^T\right) x \\ &= (C_1 \otimes I + I \otimes C_2)x \\ &= (A_2^{-1} \otimes B_1^{-1})b.\end{aligned} \tag{44}$$

Since the matrices Q_1 and Q_2 are orthogonal, the inverse of their Kronecker product $Q_1 \otimes Q_2$ is trivially given by $Q_1^T \otimes Q_2^T$. Thus, solving the system (44) is reduced to solving a system of the form $T_1 \otimes I + I \otimes T_2$. Well-known solution techniques exist for this Sylvester-type system of equations, which can be solved in $\mathcal{O}(p^3)$ operations. Thus, once the approximate preconditioner $P = A_1 \otimes B_1 + A_2 \otimes B_2$ has been computed, solving linear systems of the form $Px = b$ can be performed in linear time per degree of freedom.

4.2.1 Efficient computation of $\widetilde{A}v$ and $\widetilde{A}^T v$ in two dimensions

One of the key operations in efficiently computing the approximate Kronecker-product preconditioner is the fast, shuffled matrix-vector product operation used in the Lanczos algorithm. Since our algorithm avoids the explicit construction and evaluation of the entries of the matrix A, we present a matrix-free algorithm to compute the shuffled product in $\mathcal{O}(p^3)$ time. Setting $A = M - \Delta t J$, we apply Algorithm 3 to compute the shuffled product $\widetilde{A}u$ for a given vector u. We first write

$$\widetilde{A} = \widetilde{M} - \Delta t(\widetilde{J_{\mathrm{v}}} + \widetilde{J}_{\mathrm{f}}), \tag{45}$$

where J_{v} and J_{f} are the volume and face contributions to the Jacobian matrix, respectively. We first demonstrate the computation of the product $\widetilde{M}v$. Recall that the entries of M are given by

$$M_{ij,k\ell} = \int_K \tilde{\Phi}_{ijk}(x, y, z)\, dx. \tag{46}$$

Then, Algorithm 3 allows us to write $u = \widetilde{A}v$

$$u_{:i} = \sum_j (M_{::,ji})^T v_{:j}. \tag{47}$$

Writing out the matrix-vector product explicitly, and expanding the integral in (46) as a sum over quadrature nodes, we have

$$u_{ki} = \sum_j \sum_\ell \sum_\alpha \sum_\beta \phi_\ell(x_\alpha)\phi_k(x_\beta)\phi_j(x_\alpha)\phi_i(x_\beta)|\det(J_T(x_\alpha, x_\beta))|w_\alpha w_\beta v_{\ell j}. \tag{48}$$

This sum can be factorized as

$$u_{ki} = \sum_\beta w_\beta \phi_k(x_\beta)\phi_i(x_\beta) \sum_\alpha w_\alpha |\det(J_T(x_\alpha, x_\beta))| \sum_j \phi_j(x_\alpha) \sum_\ell \phi_\ell(x_\alpha) v_{\ell j}, \tag{49}$$

where we notice that each summation in this expression involves no more than two free indices, and therefore the expression can be computed in $\mathcal{O}(p^3)$ time.

Following the same procedure, and recalling the representation for J_{v} and J_{f} given in (27), we can evaluate the shuffled product $u = \widetilde{J_{\mathrm{v}}} v$ as

$$u_{ki} = \sum_{\beta} w_\beta \phi_k(x_\beta)\phi_i(x_\beta) \sum_{\alpha} w_\alpha \frac{\partial F_1}{\partial u}(x_\alpha, x_\beta) \sum_j \phi_j(x_\alpha) \sum_\ell \phi'_\ell(x_\alpha) v_{\ell j} + \sum_{\beta} w_\beta \phi'_k(x_\beta)\phi_i(x_\beta) \sum_{\alpha} w_\alpha \frac{\partial F_2}{\partial u}(x_\alpha, x_\beta) \sum_j \phi_j(x_\alpha) \sum_\ell \phi_\ell(x_\alpha) v_{\ell j}. \quad (50)$$

Finally, we consider the face integral terms, and write out the factorized form of the shuffled product $u = \widetilde{J_{\mathrm{f}}} v$, which takes the form

$$u_{ki} = - \sum_{e \in \partial K} \sum_{\alpha} w_\alpha \phi_i(y^e_\alpha)\phi_k(y^e_\alpha) \frac{\partial \widehat{F}}{\partial u^-}(x^e_\alpha, y^e_\alpha) \sum_j \phi_j(x^e_\alpha) \sum_\ell \phi_\ell(x^e_\alpha) v_{\ell j}. \quad (51)$$

We further remark that many of the terms in this sum can be eliminated by using the fact that many of the basis functions are identically zero along a given face e of the element K.

Computation of the transpose of the shuffled product $\widetilde{A^T} v$ can be performed using a very similar matrix-free approach, following the framework of Algorithm 4. These two procedures allow for the computation of steps 5 and 10 in the Lanczos algorithm in $\mathcal{O}(p^3)$ time.

4.3 Three spatial dimensions

In the case of three spatial dimensions, it would be natural to consider a preconditioner matrix P of the form

$$A \approx P = A_1 \otimes B_1 \otimes C_1 + A_2 \otimes B_2 \otimes C_2 + A_3 \otimes B_3 \otimes C_3. \quad (52)$$

Unfortunately, it is not readily apparent how to solve a general system of the form (52). Therefore, we instead look for a preconditioner that has the simplified form

$$A \approx P = A_1 \otimes B_1 \otimes C_1 + A_1 \otimes B_2 \otimes C_2, \quad (53)$$

where we emphasize that the same matrix A_1 appears in both terms on the right-hand side. This has the advantage that the system $Px = b$ can be transformed by multiplying on the left by $A_1^{-1} \otimes B_2^{-1} \otimes C_1^{-1}$ to obtain

$$(I \otimes B_2^{-1} B_1 \otimes I + I \otimes I \otimes C_1^{-1} C_2) x = (A_1^{-1} \otimes B_2^{-1} \otimes C_1^{-1}) b. \quad (54)$$

Applying the same technique as in the two-dimensional case allows us to simultaneously quasi-triangularize both terms on the left-hand side, which then results in a system of the form

$$(I \otimes Q_1 \otimes Q_2)(I \otimes T_1 \otimes I + I \otimes I \otimes T_2)(I \otimes Q_1^T \otimes Q_2^T) x = (A_1^{-1} \otimes B_2^{-1} \otimes C_1^{-1}) b, \quad (55)$$

which, as in the case of the two-dimensional system, is a Sylvester-type system that can be efficiently solved in $\mathcal{O}(p^3)$ time (constant time in p per degree of freedom).

4.3.1 Forming the three-dimensional preconditioner

We now address how to generate an effective preconditioner of the form (53) using the KSVD. First, we recall that the element Jacobian will be a $(p+1)^3 \times (p+1)^3$ matrix. We wish to approximate this matrix by a Kronecker product $A_1 \otimes D_1$, where $A_1 \in \mathbb{R}^{(p+1)\times(p+1)}$ and $D_1 \in \mathbb{R}^{(p+1)^2\times(p+1)^2}$. We find such matrices A_1 and D_1 by finding the largest singular value and corresponding singular vectors of the permuted matrix $\widetilde{A}$, obtaining

$$A \approx A_1 \otimes D_1. \tag{56}$$

In order to find the singular values using the Lanczos algorithm, we must compute the matrix-vector product $\widetilde{A}x$ and $\widetilde{A}^T x$. In the following section, we describe how to perform Algorithms 3 and 4 efficiently by taking advantage of the tensor-product structure of the Jacobian. Once the matrices A_1 and D_1 have been obtained, we can then repeat the KSVD process to find the best two-term approximation

$$D_1 \approx B_1 \otimes C_1 + B_2 \otimes C_2. \tag{57}$$

This too involves the Lanczos algorithm, but since the matrix D_1 has dimensions $(p+1)^2 \times (p+1)^2$, computing the permuted products $\widetilde{D_1}x$ and $\widetilde{D_1}^T x$ using standard dense linear algebra requires $\mathcal{O}(p^4)$ operations, and thus is linear in p per degree of freedom. Combining (56) and (57), we obtain an approximation of the form

$$A \approx A_1 \otimes B_1 \otimes C_1 + A_1 \otimes B_2 \otimes C_2 \tag{58}$$

as desired.

4.3.2 Efficient computation of $\widetilde{A}v$ and $\widetilde{A}^T v$ in three dimensions

As in Section 4.2.1, we describe the matrix-free procedure for computing the shuffled matrix-vector products $\widetilde{A}v$ and $\widetilde{A}^T v$. The general approach to this method is the same as in the two-dimensional case, but there are several key differences that increase the complexity of this problem. First, we recall that since our approximation takes the form $A \approx A_1 \otimes D_1$ where A_1 is $(p+1) \times (p+1)$ and D is $(p+1)^2 \times (p+1)^2$, the matrix $\tilde{A}$ is rectangular, with dimensions $(p+1)^2 \times (p+1)^4$. The algorithm we describe has linear complexity per degree of freedom of the vector v, which results in $\mathcal{O}(p^5)$ operations for the product $\widetilde{A}v$, unfortunately not meeting our overall goal of linear time per degree of freedom in the solution vector.

As before, we decompose the matrix $\widetilde{A} = \widetilde{M} - \Delta t(\widetilde{J_\mathrm{v}} + \widetilde{J}_\mathrm{f})$. First, we describe the method for the mass matrix. Recall that the entries of M are given by

$$M_{ijk,\ell mn} = \sum_{\alpha,\beta,\gamma} w_\alpha w_\beta w_\gamma |\det(J_T(x_\alpha, x_\beta, x_\gamma))| \phi_i(x_\alpha)\phi_j(x_\beta)\phi_k(x_\gamma)\phi_\ell(x_\alpha)\phi_m(x_\beta)\phi_n(x_\gamma)\, dx. \tag{59}$$

Then, following Algorithm 3, we write $u = \widetilde{M}v$, where u is a vector of length $(p+1)^2$ and v is a vector of length $(p+1)^4$,

$$u_{:i} = \sum_j \sum_k (M_{:::,jki})^T v_{::jk}. \tag{60}$$

Following the same sum factorization procedure as in the two-dimensional case, we can write

$$u_{\ell i} = \sum_{\gamma} \phi_i(x_\gamma)\phi_\ell(x_\gamma) \sum_{\beta}\sum_{\alpha} w_\alpha w_\beta w_\gamma |\det(J_T(x_\alpha, x_\beta, x_\gamma))| \sum_k \phi_k(x_\beta) \sum_j \phi_j(x_\alpha) \sum_n \phi_n(x_\beta) \sum_m \phi_m(x_\alpha) v_{mnjk}. \quad (61)$$

The $\mathcal{O}(p^5)$ complexity is clear from this form, as, for example, the right-most summation has four free indices. The shuffled product with the Jacobian of the volume integral takes a similar form, where $u = \widetilde{J_{\mathrm{v}}} v$ can be written as

$$\begin{aligned} u_{\ell i} = \sum_{\gamma} \phi_i(x_\gamma)\phi_\ell(x_\gamma) \sum_{\beta}\sum_{\alpha} w_\alpha w_\beta w_\gamma & \left(\frac{\partial F_1}{\partial u}(x_\alpha, x_\beta, x_\gamma) \sum_k \phi_k(x_\beta) \sum_j \phi_j(x_\alpha) \sum_n \phi_n(x_\beta) \sum_m \phi'_m(x_\alpha) v_{mnjk} \right. \\ & \left. + \frac{\partial F_2}{\partial u}(x_\alpha, x_\beta, x_\gamma) \sum_k \phi_k(x_\beta) \sum_j \phi_j(x_\alpha) \sum_n \phi'_n(x_\beta) \sum_m \phi_m(x_\alpha) v_{mnjk} \right) \\ & + \sum_{\gamma} \phi_i(x_\gamma)\phi'_\ell(x_\gamma) \sum_{\beta}\sum_{\alpha} w_\alpha w_\beta w_\gamma \frac{\partial F_3}{\partial u}(x_\alpha, x_\beta, x_\gamma) \\ & \sum_k \phi_k(x_\beta) \sum_j \phi_j(x_\alpha) \sum_n \phi_n(x_\beta) \sum_m \phi_m(x_\alpha) v_{mnjk}. \quad (62) \end{aligned}$$

Finally, we write the product corresponding to the face integral Jacobian, $u = \widetilde{J}_{\mathrm{f}} v$, as

$$u_{\ell i} = \sum_{e \in \partial K} \sum_{\beta}\sum_{\alpha} w_\alpha w_\beta \phi_i(z^e_{\alpha\beta})\phi_\ell(z^e_{\alpha\beta}) \frac{\partial \widehat{F}}{\partial u^-}(x^e_{\alpha\beta}, y^e_{\alpha\beta}, z^e_{\alpha\beta}) \sum_k \phi_k(y^e_{\alpha\beta}) \sum_j \phi_j(x^e_{\alpha\beta}) \sum_n \phi_n(y^e_{\alpha\beta}) \sum_m \phi_m(x^e_{\alpha\beta}) v_{mnjk}, \quad (63)$$

where $(x^e_{\alpha\beta}, y^e_{\alpha\beta}, z^e_{\alpha\beta})$ represents the coordinates of the quadrature nodes on the face e of element K indexed by (α, β). We remark that for each face e, two of the indices in the above expression can be eliminated. This simplification depends on the orientation of the face, and therefore we leave the full expression for the sake of generality.

4.4 Algorithm overview

Here we describe the overall algorithms used to form and apply the tensor product preconditioner. We present the algorithm for both the cases of two and three spatial dimensions. Forming the preconditioner requires the Lanczos SVD, given by Algorithm 2, and the two permuted matrix-vector multiplication kernels, given by Algorithms 3 and 4, and described in the preceding section. Computational complexities are indicated for each step of the algorithm. We note that in the 2D case, we obtain an overall complexity of $\mathcal{O}(p^{d+1})$. In the 3D case, all the operations have complexity at most $\mathcal{O}(p^{d+1})$, except the first Lanczos SVD, which requires $\mathcal{O}(p^5)$ operations.

Algorithm 5 Form 2D preconditioner $A \approx P = A_1 \otimes B_2 + A_2 \otimes B_2$

1: Compute $A \approx A_1 \otimes B_1 + A_2 \otimes B_2$ using Lanczos iteration and matrix-free products $\widetilde{A}x$ and $\widetilde{A}^T x$ ▷ *Complexity:* $\mathcal{O}(p^3)$
2: Precompute LU factorizations of A_2 and B_1 ▷ *Complexity:* $\mathcal{O}(p^3)$
3: Precompute Schur factorizations $Q_1 T_1 Q_1^T, Q_2 T_2 Q_2^T$ of $A_2^{-1} A_1$ and $B_1^{-1} b_2$, respectively ▷ *Complexity:* $\mathcal{O}(p^3)$

Algorithm 6 Apply 2D preconditioner to solve $Px = b$

1: $\tilde{b} \leftarrow A_2^{-1} \otimes B_1^{-1} b$ ▷ *Complexity:* $\mathcal{O}(p^3)$
2: Solve the Sylvester system $(T_1 \otimes I + I \otimes T_2)\tilde{x} = (Q_1^T \otimes Q_2^T)\tilde{b}$ ▷ *Complexity:* $\mathcal{O}(p^3)$
3: $x \leftarrow (Q_1 \otimes Q_2)\tilde{x}$ ▷ *Complexity:* $\mathcal{O}(p^3)$

Algorithm 7 Form 3D preconditioner $A \approx P = A_1 \otimes B_1 \otimes C_1 + A_1 \otimes B_2 \otimes C_2$

1: Compute $A \approx A_1 \otimes D_1$ using Lanczos iteration and matrix-free products $\widetilde{A}x$ and $\widetilde{A}^T x$ ▷ *Complexity:* $\mathcal{O}(p^5)$
2: Compute $D_1 \approx B_1 \otimes C_1 + B_2 \otimes C_2$ using Lanczos iteration and dense permuted products ▷ *Complexity:* $\mathcal{O}(p^4)$
3: Precompute LU factorizations of A_1, B_2, and C_1 ▷ *Complexity:* $\mathcal{O}(p^3)$
4: Precompute Schur factorizations $Q_1 T_1 Q_1^T, Q_2 T_2 Q_2^T$ of $B_2^{-1} B_1$ and $C_1^{-1} C_2$, respectively ▷ *Complexity:* $\mathcal{O}(p^3)$

Algorithm 8 Apply 3D preconditioner to solve $Px = b$

1: $\tilde{b} \leftarrow A_1^{-1} \otimes B_2^{-1} \otimes C_1^{-1} b$ ▷ *Complexity:* $\mathcal{O}(p^4)$
2: Solve the Sylvester system $(I \otimes T_1 \otimes I + I \otimes I \otimes T_2)\tilde{x} = (I \otimes Q_1^T \otimes Q_2^T)\tilde{b}$ ▷ *Complexity:* $\mathcal{O}(p^4)$
3: $x \leftarrow (I \otimes Q_1 \otimes Q_2)\tilde{x}$ ▷ *Complexity:* $\mathcal{O}(p^4)$

5 Numerical results

5.1 2D linear advection equation

The simplest example we consider is that of the two-dimensional scalar advection equation, given by

$$u_t + \nabla \cdot (\alpha, \beta)\, u = 0, \tag{64}$$

where (α, β) is a space-dependent velocity field. Because of the particularly simple structure of this equation, it is possible to see how the approximate Kronecker preconditioner, given by $A_1 \otimes B_1 + A_2 \otimes B_2$ relates to the true diagonal blocks A of the matrix $M - \Delta t J$. In this case, the properties of the velocity field (α, β) can determine how well the discontinuous Galerkin discretization can be approximated by a tensor-product structure.

Neglecting for now the face integral terms, the diagonal blocks of $M - \Delta t J$ can be written as

$$\left(G^T W \otimes G^T W\right) J_T \left(G \otimes G\right) - \Delta t \left(\left(G^T W \otimes D^T W\right) F_1 + \left(D^T W \otimes G^T W\right) F_2\right) \left(G \otimes G\right), \quad (65)$$

where J_T, F_1, and F_2 are $\mu^2 \times \mu^2$ diagonal matrices. If the matrices J_T, F_1, and F_2 additionally posses a Kronecker-product structure, then it is possible to rewrite (65) exactly in the form $A_1 \otimes B_1 + A_2 \otimes B_2$. The Kronecker structure of J_T is determined by the geometry of the mesh, and the structure of F_1 and F_2 is determined by the form of the velocity field.

For example, we first suppose that the mesh is a cartesian grid with grid size h, and thus the Jacobian determinant of the transformation map is equal to h^2. Hence, J_T is equal to h^2 times the identity matrix. If we further suppose that the velocity field is separable, in the sense that, each component depends only on the corresponding spatial variable, *i.e.* $\alpha(x, y) = \alpha(x)$, $\beta(x, y) = \beta(y)$, then the flux derivatives can be written as $F_1 = I \otimes F_{1_x}$, and $F_2 = F_{2_y} \otimes I$. Therefore, we can rewrite (65) in the form

$$\left(h^2 G^T W G - \Delta t D^T W F_{2_y} G\right) \otimes G^T W G - \Delta t \left(G^T W G \otimes D^T W F_{1_x} G\right), \quad (66)$$

and we see that the diagonal blocks of $M - \Delta t J$ are exactly representable by our Kronecker-product approximation.

If, on the other hand, we allow straight-sided, non-cartesian meshes, then the transformation mapping is a bilinear function, and its Jacobian determinant is a linear function in the variables x and y. Thus, $J_T = J_{T_y} \otimes I + I \otimes J_{T_x}$. If the velocity field is constant in space, then we obtain the following representation of the diagonal blocks

$$G^T W G \otimes \left(G^T W J_{T_x} G - \Delta t \alpha D^T W G\right) + \left(G^T W J_{T_y} G - \Delta t \beta D^T W G\right) \otimes G^T W G, \quad (67)$$

and we see that our Kronecker-product approximation is again exact.

If, in contrast to the previous two cases, the transformation mapping is given by a higher degree polynomial, or if the velocity field is not separable, then the approximate preconditioner will not yield the exact diagonal blocks. However, if the deformation of the mesh is not too large, if the velocity field is well approximated by one that is separable, or if the time step Δt is relatively small, then we expect the Kronecker product preconditioner to compare favorably with the exact block Jacobi preconditioner.

We remark that this numerical experiment is designed to highlight two main features of the Kronecker-product preconditioner. The first is that if any exact representation of the diagonal blocks in the form $A_1 \otimes B_1 + A_2 \otimes B_2$ exists, such as those given by equations (66) and (67), then the KSVD algorithm provides an *automatic*, and purely *algebraic* method to identity this decomposition. No special structure of the flux functions is required to be known a priori in order for the KSVD to exactly reproduce this tensor-product structure. Secondly, even in a case where it is impossible to write such an expression exactly, the KSVD method will identify the best possible approximation of this form. Thus, in cases where the velocity field is close to constant (*e.g.* when the mesh size is very small), or where the mesh deformation is small, we expect this approximation to be very accurate.

In order to compare the performance of these two preconditioners, we solve equation (64) on both regular and irregular meshes, with constant, separable, and non-separable velocity fields. The meshes are shown in Figure 1, and the velocity fields in Figure 2. We choose a representative time step of $\Delta t = 0.5$, and consider polynomial degrees $p = 1, 2, \ldots, 10$. In the case of the regular cartesian mesh, we expect identical performance for the exact block Jacobi and approximate

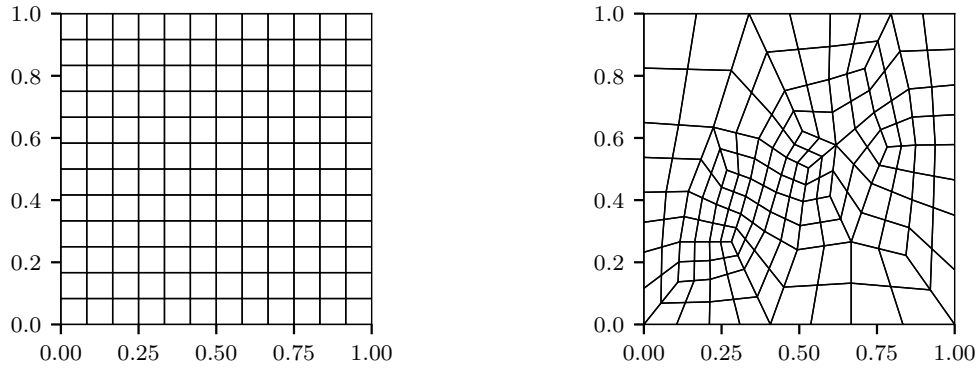

(a) Cartesian mesh for advection equation

(b) Unstructured mesh for advection equation

Figure 1: Meshes used for the advection equation

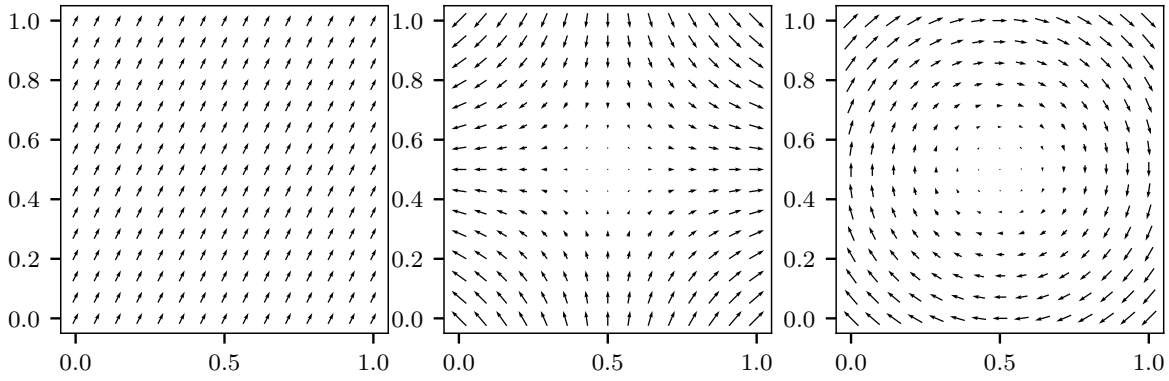

(a) $\alpha = 1, \beta = 2$

(b) $\alpha = x - \frac{1}{2}, \beta = \frac{1}{2} - y$

(c) $\alpha = y - \frac{1}{2}, \beta = \frac{1}{2} - x$

Figure 2: Velocity fields used for the advection equation

Kronecker-product preconditioners for the velocity fields shown in Figures 2a and 2b, since the diagonal blocks can be reproduced exactly. In the case of the unstructured mesh, we expect to see identical performance for the constant velocity field in Figure 2a. Indeed, the numerical results corroborate our expectations, and the number of iterations is identical between the two preconditioners in those test cases. Additionally, even in cases where the Kronecker-product approximation cannot reproduce the exact blocks, such for for velocity field 2c or 2b on the unstructured mesh, the performance is, in most cases, extremely similar to that of exact block Jacobi. The number of GMRES iterations required to converge with each preconditioner is shown in Table 2. For very large values of polynomial degree p on the unstructured mesh, with non-separable velocity field, we begin to see a degradation in the performance of the Kronecker-product preconditioner.

(a) Cartesian grid

p	**(a)** J	**(a)** K	**(b)** J	**(b)** K	**(c)** J	**(c)** K
1	12	12	5	5	29	29
2	14	14	7	7	29	29
3	13	13	7	7	29	29
4	14	14	7	7	29	29
5	13	13	7	7	29	30
6	17	17	7	7	29	31
7	13	13	7	7	30	28
8	14	14	7	7	30	29
9	12	12	7	7	27	30
10	14	14	7	7	27	30

(b) Unstructured mesh

p	**(a)** J	**(a)** K	**(b)** J	**(b)** K	**(c)** J	**(c)** K
1	14	14	10	11	29	29
2	15	15	11	10	29	29
3	14	14	11	12	28	28
4	15	15	9	12	28	31
5	14	14	10	12	28	34
6	14	14	10	12	28	39
7	13	13	10	12	28	46
8	13	13	12	13	28	53
9	13	13	12	14	28	62
10	13	13	12	15	28	69

Table 2: Number of GMRES iterations for Jacobi (J) and KSVD (K) preconditioners, advection equation on cartesian and unstructured grids, velocity fields (a), (b), and (c) from Figure 2.

5.2 2D Euler vortex

In this example, we consider the compressible Euler equations of gas dynamics in two dimensions, given in conservative form by

$$\partial_t \boldsymbol{u} + \nabla \cdot F(\boldsymbol{u}) = 0, \tag{68}$$

where

$$\boldsymbol{u} = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{pmatrix}, \qquad F_1(\boldsymbol{u}) = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho H u \end{pmatrix}, \qquad F_2(\boldsymbol{u}) = \begin{pmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho H v \end{pmatrix}, \tag{69}$$

where ρ is the density, $\boldsymbol{v} = (u, v)$ is the fluid velocity, p is the pressure, and E is the specific energy. The total enthalpy H is given by

$$H = E + \frac{p}{\rho}, \tag{70}$$

and the pressure is determined by the equation of state

$$p = (\gamma - 1)\rho \left(E - \frac{1}{2}(u^2 + v^2) \right), \tag{71}$$

where $\gamma = c_p/c_v$ is the ratio of specific heat capacities at constant pressure and constant volume, taken to be 1.4.

We consider the model problem of an unsteady compressible vortex in a rectangular domain [28]. The domain is taken to be a 20×15 rectangle and the vortex is initially centered at $(x_0, y_0) = (5, 5)$.

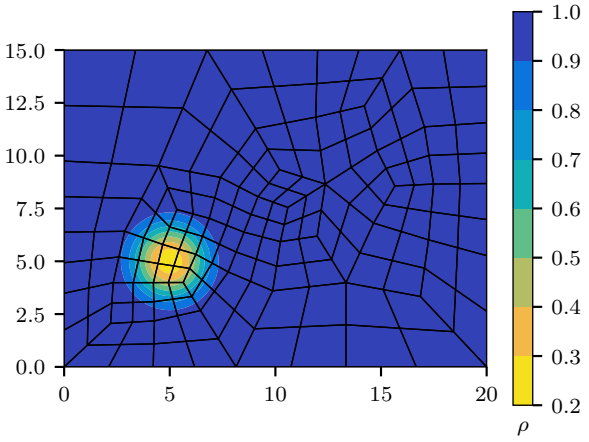

Figure 3: Initial conditions (density) for Euler vortex on unstructured mesh.

The vortex is moving with the free-stream at an angle of θ. The exact solution is given by

$$u = u_\infty \left(\cos(\theta) - \frac{\epsilon((y - y_0) - \overline{v}t)}{2\pi r_c} \exp\left(\frac{f(x,y,t)}{2} \right) \right), \tag{72}$$

$$u = u_\infty \left(\sin(\theta) - \frac{\epsilon((x - x_0) - \overline{u}t)}{2\pi r_c} \exp\left(\frac{f(x,y,t)}{2} \right) \right), \tag{73}$$

$$\rho = \rho_\infty \left(1 - \frac{\epsilon^2(\gamma - 1)M_\infty^2}{8\pi^2} \exp(f(x,y,t)) \right)^{\frac{1}{\gamma - 1}}, \tag{74}$$

$$p = p_\infty \left(1 - \frac{\epsilon^2(\gamma - 1)M_\infty^2}{8\pi^2} \exp(f(x,y,t)) \right)^{\frac{\gamma}{\gamma - 1}}, \tag{75}$$

where $f(x,y,t) = (1 - ((x - x_0) - \overline{u}t)^2 - ((y - y_0) - \overline{v}t)^2)/r_c^2$, M_∞ is the Mach number, u_∞, ρ_∞, and p_∞ are the free-stream velocity, density, and pressure, respectively. The free-stream velocity is given by $(\overline{u}, \overline{v}) = u_\infty(\cos(\theta), \sin(\theta))$. The strength of the vortex is given by ϵ, and its size is r_c. We choose the parameters to be $M_\infty = 0.5$, $u_\infty = 1$, $\theta = \arctan(1/2)$, $\epsilon = 0.3$, and $r_c = 1.5$.

As in the case of the linear advection equation, we consider both a regular $n_x \times n_y$ cartesian grid, and an unstructured mesh. The unstructured mesh is obtained by scaling the mesh in Figure 1b by 20 in the x-direction and 15 in the y-direction. Density contours of the initial conditions are shown in Figure 3. As opposed to the scalar advection equation, the solution to the Euler equations consists of multiple components. Thus, the blocks of the Jacobian matrix can be considered to be of size $n_c(p+1)^2 \times n_c(p+1)^2$, where n_c is the number of solution components (in the case of the 2D Euler equations, $n_c = 4$). The exact block Jacobi preconditioner computes the inverses of these large blocks. The approximate Kronecker-product preconditioner find optimal approximation of the form $A_1 \otimes B_1 + A_2 \otimes B_2$, where A_1 and A_2 are $n_c(p+1) \times n_c(p+1)$ matrices, and B_1 and B_2 are $(p+1) \times (p+1)$ matrices.

We choose two representative time steps of $\Delta t = 0.1$ and $\Delta t = 0.01$, and compute the average number of GMRES iterations per linear solve required to perform one backward Euler time step. We choose the polynomial degree $p = 3, 4, \ldots, 15$, and consider cartesian and unstructured meshes,

both with 160 quadrilateral elements. We present the results in Table 4. Very similar results are observed for the structured and unstructured results. We note that for the smaller time step, the approximate Kronecker-product preconditioner requires a very similar number of iterations when compared with the exact block Jacobi preconditioner, even for high polynomial degree p. For the larger time step, the number of iterations required for the KSVD preconditioner increases with p at a faster rate when compared with the block Jacobi preconditioner, suggesting that the Kronecker-product preconditioner is most effective for moderate time steps Δt.

5.2.1 Performance comparison

In this section we compare the runtime performance of the Kronecker-product preconditioner with the exact block Jacobi preconditioner. Although we have observed that for large time steps Δt or polynomial degrees p, the KSVD preconditioner requires more iterations to converge, it is also possible to compute and apply this preconditioner much more efficiently. Here, we compare the wall-clock time required to compute and apply the preconditioner, according to Algorithms 5 and 6, respectively, for $p = 3, 4, \ldots, 15$. The block Jacobi preconditioner is computed by first assembling the diagonal block of the Jacobian matrix using the sum-factorized form of expression (27), and then computing its LU factorization. The wall-clock times for these operations are shown in Figure 4a. We remark that we observe the expected asymptotic computational complexities for each of these operations, where forming the Jacobi preconditioner requires $\mathcal{O}(p^6)$ operations, and applying the Jacobi preconditioner requires $\mathcal{O}(p^4)$ operations. Both forming and applying the approximate Kronecker-product preconditioner can be done in $\mathcal{O}(p^3)$ time. The total runtime observed per backward Euler step for $\Delta t = 0.1$ and $\Delta t = 0.01$ is shown in Figure 4b. We see that for $p \geq 7$ and $\Delta t = 0.01$, the Kronecker-product preconditioner results in overall faster runtime, while for

(a) Cartesian grid

	$\Delta t = 0.01$		$\Delta t = 0.1$	
p	J	K	J	K
3	5	6	11	18
4	6	7	12	23
5	6	8	13	30
6	7	9	15	38
7	7	10	17	47
8	8	11	18	59
9	8	13	20	71
10	9	15	21	88
11	9	17	23	103
12	10	19	25	121
13	11	20	24	123
14	11	23	25	157
15	12	25	26	196

(b) Unstructured mesh

	$\Delta t = 0.01$		$\Delta t = 0.1$	
p	J	K	J	K
3	6	7	12	19
4	6	7	14	26
5	7	9	16	32
6	7	10	17	42
7	8	11	18	51
8	8	12	20	64
9	9	13	21	74
10	9	15	23	90
11	9	17	24	110
12	10	19	25	125
13	10	21	25	142
14	11	24	26	164
15	11	26	27	245

Table 4: Number of GMRES iterations for Jacobi (J) and KSVD (K) preconditioners, Euler equations on cartesian and unstructured grids, with $\Delta t = 0.1$ and $\Delta t = 0.01$.

$\Delta t = 0.1$, because of the large number of iterations required per solve, the Jacobi preconditioner results in overall faster performance.

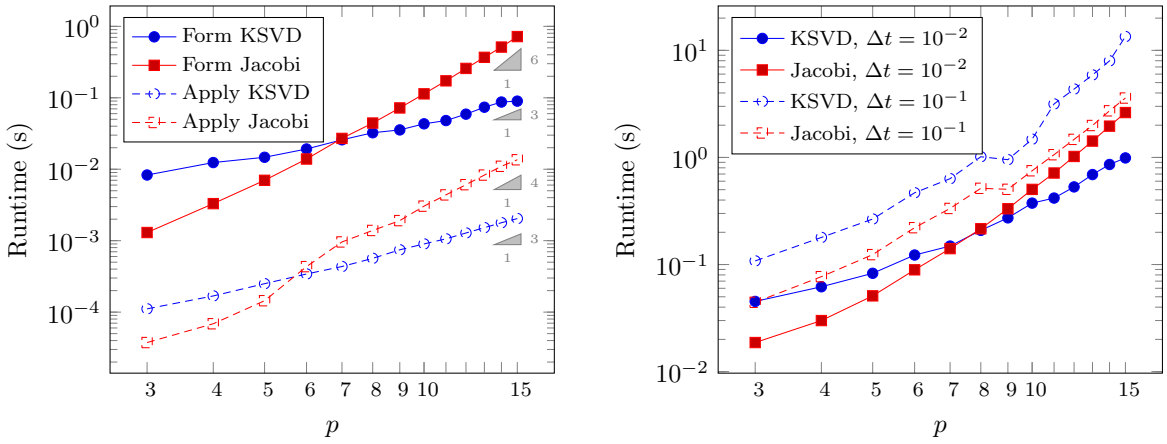

(a) Wall-clock time required to form and apply the KSVD and Jacobi preconditioners

(b) Total wall-clock time required per backward Euler step

Figure 4: Runtime performance comparison of Kronecker-product preconditioner with exact block Jacobi preconditioner for 2D Euler equations.

5.3 2D Kelvin-Helmholtz instability

For a more sophisticated test case, we consider a two-dimensional Kelvin-Helmholtz instability. This important fluid instability occurs in shear flows of fluids with different densities. The domain is taken to be the periodic unit square $[0, 1]^2$. We define the function

$$f(x) = \frac{1}{4}(\text{erf}(\alpha(x - 0.25)) + 1)(\text{erf}(\alpha(0.75 - x)) + 1),$$

where $\alpha = 100$, as a smooth approximation to the discontinuous characteristic function

$$\chi(x) = \begin{cases} 1, & 0.25 \leq x \leq 0.75 \\ 0, & \text{otherwise.} \end{cases}$$

Following [25], we define the initial conditions by

$$\rho(x, y) = f(y) + 1, \qquad u(x, y) = f(y) - 1/2, \qquad p(x, y) = 2.5,$$

where the vertical velocity is given by

$$v(x, y) = \frac{1}{10}\sin(4\pi x)\left(\exp\left(-\frac{(y - 0.25)^2}{2\sigma^2}\right) + \exp\left(-\frac{(y - 0.75)^2}{2\sigma^2}\right)\right).$$

Thus, the fluid density is equal to 2 inside the strip $y \in [0.25, 0.75]$, and 1 outside the strip. The fluid is moving to the right with horizontal velocity 0.5 inside the strip, and is moving to the left

with equal speed outside of the strip. A small perturbation in the vertical velocity, localized around the discontinuity, determines the large-scale behavior of the instability. The initial conditions are shown in Figure 5.

We use a 128×128 cartesian grid, with polynomial bases of degree 3, 7, and 10. For 10th degree polynomials, the total number of degrees of freedom is 7,929,856. The Euler equations are integrated for 1.5 s using a fourth-order explicit method with a time step of $\Delta t = 2.5 \times 10^{-5}$ on the NERSC Edison supercomputer, using 480 cores. At this point, the solution has developed sophisticated large- and small-scale features, as shown in Figure 6.

We then linearize the Euler equations around this solution in order to test the preconditioner performance. Because of the varied scale of the features in this solution, we believe that the resulting linearization is a representative of the DG systems we are interested in solving. Using this solution, we then solve one backward Euler step using both the Jacobi and approximate Kronecker-product preconditioners. For the implicit solve, we choose a range of time steps, from the explicit step size of $\Delta t = 2.5 \times 10^{-5}$, to a larger step of $\Delta t = 10^{-3}$. The number of GMRES iterations required per linear solve are show in Table 6. We observe that for the explicit-scale time step, the exact block Jacobi and approximation Kronecker-product preconditioner exhibit very similar performance for all choices of p. For the largest time step, $\Delta t = 10^{-3}$, the Kronecker-product preconditioner required about twice as many iterations for $p = 3$, three times as many for $p = 7$, and four times as many iterations for $p = 10$.

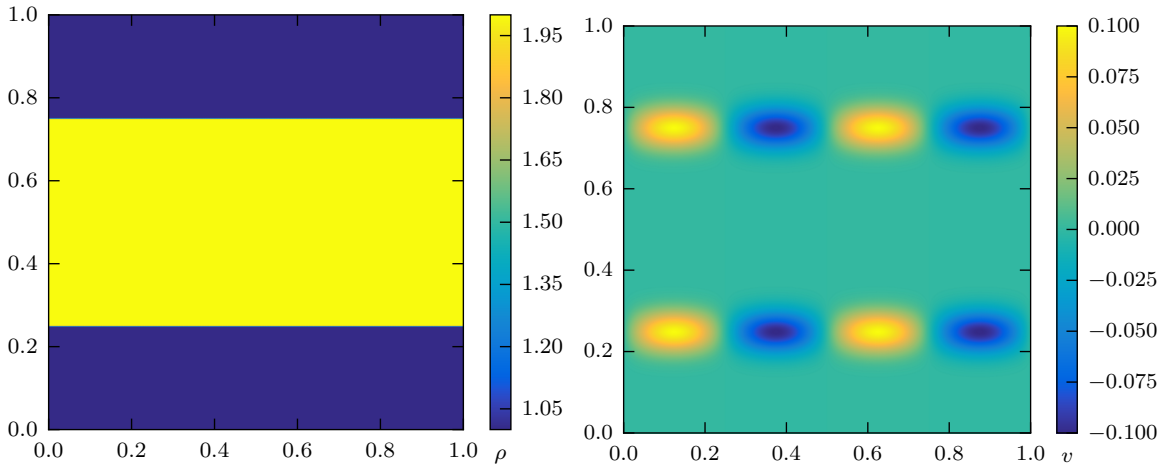

(a) Density

(b) Vertical velocity

Figure 5: Initial conditions for the Kelvin-Helmholtz instability.

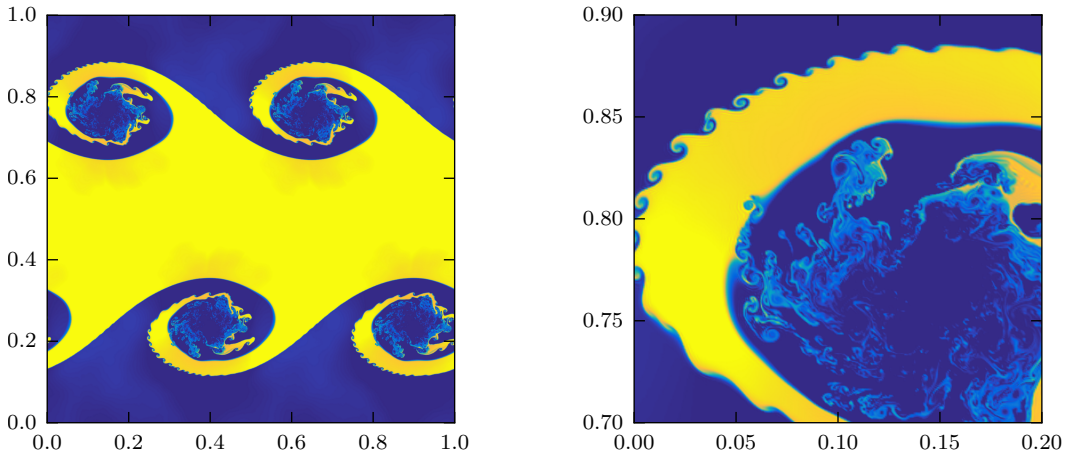

Figure 6: Solution (density) of Kelvin-Helmholtz instability at $t = 1.5$ s, with zoomed-in subregion to show small-scale features.

$p = 3$

Δt	J	K
2.5×10^{-5}	4	4
5.0×10^{-5}	5	5
1.0×10^{-4}	6	7
2.5×10^{-4}	8	10
5.0×10^{-4}	10	14
1.0×10^{-3}	13	22

$p = 7$

Δt	J	K
2.5×10^{-5}	5	6
5.0×10^{-5}	6	8
1.0×10^{-4}	8	12
2.5×10^{-4}	12	20
5.0×10^{-4}	15	35
1.0×10^{-3}	21	62

$p = 10$

Δt	J	K
2.5×10^{-5}	6	8
5.0×10^{-5}	8	11
1.0×10^{-4}	10	16
2.5×10^{-4}	14	31
5.0×10^{-4}	19	55
1.0×10^{-3}	24	106

Table 6: Number of GMRES iterations for Jacobi (J) and KSVD (K) preconditioners, Euler equations for 2D Kelvin-Helmholtz instability.

5.4 3D periodic Euler

In this example, we provide a test case for the three-dimensional preconditioner. We consider the cube $[0, 2]^3$ with periodic boundary conditions, and use the initial conditions from [10], given by

$$\rho = 1 + 0.2 \sin(\pi(x + y + z)), \tag{76}$$

$$u = 1, \quad v = -1/2, \quad w = 1, \tag{77}$$

$$p = 1. \tag{78}$$

The exact solution to the Euler equations is known analytically in this case. Velocity and pressure remain constant in time, and the density at time t is given by

$$\rho = 1 + 0.2 \sin(\pi(x + y + z - t(u + v + w))). \tag{79}$$

The initial conditions are shown in Figure 7. The mesh is taken to be a regular $6 \times 6 \times 6$ hexahedral grid, and we consider polynomial degrees of $p = 1, 2, \ldots, 12$. We choose a representative time step

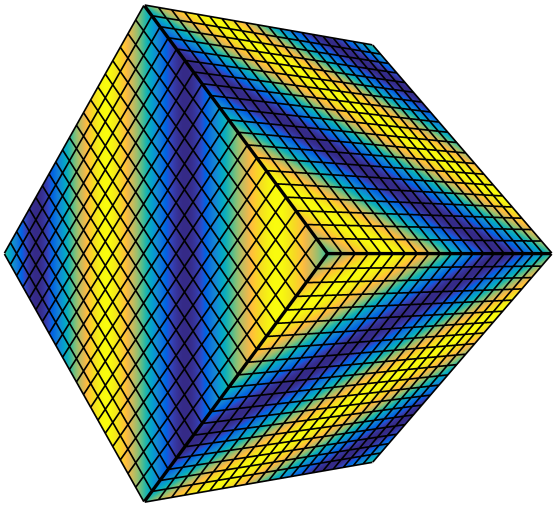

Figure 7: Initial conditions (density) for smooth 3D Euler test case.

of $\Delta t = 2.5 \times 10^{-3}$. This time step can be used for explicit methods with low-degree polynomials, but because of the p-dependency of the CFL condition, we observe that for $p \geq 9$, explicit methods become unstable, motivating the use of implicit methods. In order to compare the efficiency of the approximate Kronecker-product preconditioner, we compute one backward Euler step and compare the number of GMRES iterations required per linear solve using the exact block Jacobi preconditioner and the KSVD preconditioner.

In this test case, we also consider two variations each of these preconditioners. Since the solution to the Euler equations consists of five components, we can consider the diagonal blocks of the Jacobian to either be large $5(p+1)^3 \times 5(p+1)^3$ blocks coupling all solution components, or as smaller $(p+1)^3 \times (p+1)^3$ blocks, which do not couple the solution components. We expect that using the smaller blocks will require more GMRES iterations per linear solve, because each block captures less information. The larger blocks, on the other hand, are much more computationally expensive to compute. We present the number of iterations required to converge for each of the preconditioners in Table 8, where J and K stand for the block Jacobi and Kronecker-product preconditioners, respectively, and the subscripts "full" and "small" refer to the block size used. We observe that for small polynomial degree p, the KSVD preconditioner results in close-to-identical number of iterations, when compared with exact block Jacobi. For this test case, for p closer to 12, we observe that the number of iterations grows faster for the KSVD preconditioner than for the block Jacobi preconditioner, but the difference remains relatively small. Additionally, as expected, the small-block preconditioner requires a greater number of iterations to converge when compared with the full-block preconditioner. We note that due to the very large memory requirements, the full Jacobi preconditioner did not complete for $p = 12$.

We additionally measure the wall-clock time required per backward Euler step for each of the preconditioners, and present the results in Figure 8. From these measurements, it is possible to draw several conclusions. Firstly, as p grows, it becomes possible to observe the $\mathcal{O}(p^9)$ complexity of the exact block Jacobi preconditioner, which becomes prohibitively expensive. This is in contrast to

p	J_{full}	K_{full}	J_{small}	K_{small}
1	4	4	5	5
2	4	5	6	6
3	5	5	7	7
4	5	5	8	8
5	5	6	9	10
6	5	6	11	12
7	5	6	12	13
8	5	7	15	16
9	5	7	16	17
10	5	8	19	21
11	5	8	20	22
12	-	9	23	26

Table 8: GMRES iterations required per linear solve for three-dimensional periodic Euler test case.

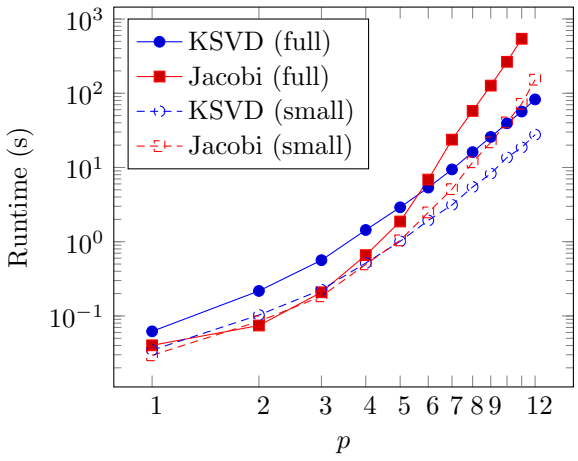

Figure 8: Wall-clock time required per backward Euler solve for three-dimensional periodic Euler test case.

the Kronecker-product preconditioner, whose $\mathcal{O}(p^5)$ complexity results in reasonable runtimes for all p considered. The full-block KSVD preconditioner results in faster runtimes starting at about $p = 5$, and the small-block preconditioner at about $p = 4$. We also see that, despite the larger number of iterations required, the small-block preconditioner results in faster overall runtime.

6 Conclusion

In this work we have developed new approximate tensor-product based preconditioners for very high-order discontinuous Galerkin methods. These preconditioners are computed using an algebraic singular value-based algorithm, and compare favorably with the traditional block Jacobi preconditioner. The computational complexity is reduced from $\mathcal{O}(p^6)$ to $\mathcal{O}(p^3)$ in two spatial dimensions and $\mathcal{O}(p^9)$ to $\mathcal{O}(p^5)$ in three spatial dimensions. Numerical results in two and three dimensions for the advection and Euler equations, using polynomial degrees up to $p = 15$, confirm the expected computational complexities, and demonstrate significant reductions in runtimes for certain test problems.

7 Acknowledgments

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The first author was supported by the Department of Defense through the National Defense Science & Engineering Graduate Fellowship Program.

References

[1] Abdalkader Baggag, Harold Atkins, and David Keyes. Parallel implementation of the discontinuous galerkin method. 1999.

[2] Philipp Birken, Gregor Gassner, Mark Haas, and Claus-Dieter Munz. Preconditioning for modal discontinuous Galerkin methods for unsteady 3D Navier-Stokes equations. *Journal of Computational Physics*, 240:20–35, may 2013.

[3] Bernardo Cockburn and Chi-Wang Shu. The Runge-Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems. *Journal of Computational Physics*, 141(2):199–224, 1998.

[4] A. Crivellini and F. Bassi. An implicit matrix-free discontinuous Galerkin solver for viscous and turbulent aerodynamic simulations. *Computers & Fluids*, 50(1):81–93, nov 2011.

[5] Laslo T. Diosady. *Domain Decomposition Preconditioners for Higher-Order Discontinuous Galerkin Discretizations*. PhD thesis, Massachusetts Institute of Technology, 2011.

[6] Laslo T. Diosady and Scott M. Murman. Tensor-product preconditioners for higher-order spacetime discontinuous Galerkin methods. *Journal of Computational Physics*, 330:296 – 318, 2017.

[7] J. A. Escobar-Vargas, P. J. Diamessis, and C. F. Van Loan. The numerical solution of the pressure Poisson equation for the incompressible Navier-Stokes equations using a quadrilateral spectral multidomain penalty method. Preprint available at https://www.cs.cornell.edu/cv/ResearchPDF/Poisson.pdf, 2011.

[8] Gene H. Golub, Franklin T. Luk, and Michael L. Overton. A block Lanczos method for computing the singular values and corresponding singular vectors of a matrix. *ACM Transactions on Mathematical Software (TOMS)*, 7(2):149–169, 1981.

[9] J. Gopalakrishnan and G. Kanschat. A multilevel discontinuous galerkin method. *Numerische Mathematik*, 95(3):527–550, 2003.

[10] Guang-Shan Jiang and Chi-Wang Shu. Efficient implementation of weighted ENO schemes. *Journal of Computational Physics*, 126(1):202 – 228, 1996.

[11] Guido Kanschat. Robust smoothers for high-order discontinuous Galerkin discretizations of advectiondiffusion problems. *Journal of Computational and Applied Mathematics*, 218(1):53 – 60, 2008.

[12] Robert Klöfkorn. Efficient matrix-free implementation of discontinuous Galerkin methods for compressible flow problems. In *Proceedings of the Conference Algoritmy*, pages 11–21, 2015.

[13] Lilia Krivodonova and Ruibin Qin. An analysis of the spectrum of the discontinuous Galerkin method. *Applied Numerical Mathematics*, 64:1–18, 2013.

[14] M. Kronbichler, S. Schoeder, C. Müller, and W. A. Wall. Comparison of implicit and explicit hybridizable discontinuous galerkin methods for the acoustic wave equation. *International Journal for Numerical Methods in Engineering*, 106(9):712–739, 2016. nme.5137.

[15] Charles F. Van Loan. The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(12):85 – 100, 2000. Numerical Analysis 2000. Vol. III: Linear Algebra.

[16] Robert E. Lynch, John R. Rice, and Donald H. Thomas. Direct solution of partial difference equations by tensor product methods. *Numerische Mathematik*, 6(1):185–199, 1964.

[17] K.-A. Mardal, T. K. Nilssen, and G. A. Staff. Order-optimal preconditioners for implicit Runge-Kutta schemes applied to parabolic PDEs. *SIAM Journal on Scientific Computing*, 29(1):361–375, 2007.

[18] Steven A. Orszag. Spectral methods for problems in complex geometries. *Journal of Computational Physics*, 37(1):70 – 92, 1980.

[19] Jaime Peraire and Per-Olof Persson. High-order discontinuous Galerkin methods for CFD. In Z. J. Wang, editor, *Adaptive High-Order Methods in Fluid Dynamics*, chapter 5, pages 119–152. World Scientific, 2011.

[20] Per-Olof Persson and Jaime Peraire. An efficient low memory implicit DG algorithm for time dependent problems. In *44th AIAA Aerospace Sciences Meeting and Exhibit*, page 113, 2006.

[21] Per-Olof Persson and Jaime Peraire. Newton-GMRES preconditioning for discontinuous Galerkin discretizations of the Navier-Stokes equations. *SIAM Journal on Scientific Computing*, 30(6):2709–2733, 2008.

[22] W. H. Reed and T. R. Hill. Triangular mesh methods for the neutron transport equation. *Los Alamos Report LA-UR-73-479*, 1973.

[23] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.

[24] Jie Shen, Tao Tang, and Li-Lian Wang. *Separable Multi-Dimensional Domains*, pages 299–366. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[25] Volker Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791, 2010.

[26] Charles F. Van Loan and Nikos Pitsianis. Approximation with Kronecker products. In *Linear Algebra for Large Scale and Real-Time Applications*, pages 293–314. Springer, 1993.

[27] Peter E. J. Vos, Spencer J. Sherwin, and Robert M. Kirby. From h to p efficiently: Implementing finite and spectral/hp element methods to achieve optimal performance for low-and high-order discretisations. *Journal of Computational Physics*, 229(13):5161–5181, 2010.

[28] Z.J. Wang, Krzysztof Fidkowski, Rémi Abgrall, Francesco Bassi, Doru Caraeni, Andrew Cary, Herman Deconinck, Ralf Hartmann, Koen Hillewaert, H.T. Huynh, Norbert Kroll, Georg May, Per-Olof Persson, Bram van Leer, and Miguel Visbal. High-order CFD methods: current status and perspective. *International Journal for Numerical Methods in Fluids*, 72(8):811–845, 2013.